

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
14 February 2002 (14.02.2002)

PCT

(10) International Publication Number
WO 02/13023 A1

(51) International Patent Classification⁷: G06F 13/00, 13/38, 17/00, H04L 12/66, H04M 1/64, 3/00, 7/00, 11/00

(21) International Application Number: PCT/US01/24878

(22) International Filing Date: 8 August 2001 (08.08.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/643,035 8 August 2000 (08.08.2000) US
09/689,222 11 October 2000 (11.10.2000) US

(71) Applicant (for all designated States except US): PROMINENCE NETWORKS, INC. [US/US]; 943 Holmdel Raod, Holmdel, NJ 07733 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): NAG, Shiddharta [US/US]; 1 Tiberon Drive, Holmdel, NJ 07733 (US). D'SOUZA, Alfred [US/US]; 42 Priscilla Drive, Lincroft, NJ 07738 (US). ALAM, Naveed [US/US]; 25 Vanderveer Road, Freehold, NJ 07726 (US). PATEL, Rakesh [US/US]; 636 Prospect Avenue, Fairview, NJ 07022 (US).

(74) Agents: MALLIE, Michael, J. et al.; Blakely, Sokoloff, Taylor & Zafman LLP, 7th floor, 12400 Wilshire Boulevard, Los Angeles, CA 90025 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, IU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.


(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

[Continued on next page]

(54) Title: MULTIPLEXING SEVERAL INDIVIDUAL APPLICATION SESSIONS OVER A PRE-ALLOCATED RESERVATION PROTOCOL SESSION IN A VOICE OVER INTERNET PROTOCOL (VOIP)



WO 02/13023 A1

(57) Abstract: Apparatus and methods are provided for multiplexing application flows over a preallocated bandwidth reservation protocol session. Additionally, a graphical user interface (GUI) is disclosed that allows a user to identify routers, communities, residents and media aggregation managers (110, 120) existing on a network. According to one embodiment, a pre-allocated reservation protocol session, such as an RSVP session, is shared by one or more application sessions. The reservation protocol session is pre-allocated over a path between a first network device (111, 112) associated with a first user community and a second network device (121, 122) associated with a second user community based upon an estimated usage of the path for application sessions between users of the first and second user communities. Subsequently, the one or more application sessions are dynamically aggregated by multiplexing application flows associated with the one or more individual application sessions onto the pre-allocated reservation protocol session at the first network device (111, 112) and demultiplexing at the second network device (121, 122).



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

MULTIPLEXING SEVERAL INDIVIDUAL APPLICATION SESSIONS OVER A PRE-ALLOCATED RESERVATION PROTOCOL SESSION IN A VOICE OVER INTERNET PROTOCOL (VOIP)

This application claims the benefit of U.S. Patent Application No. 09/634,035, filed August 8, 2000, titled "Multiplexing Several Individual Application Sessions over a Pre-Allocated Reservation Protocol Session", and U.S. Patent Application No. 09/689,222, filed October 11, 2000, titled "Graphical User Interface (GUI) For Administering A Voice Over Internet Protocol (voIP) Network Implementing Media Aggregation Managers."

COPYRIGHT NOTICE

Contained herein is material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction of the patent disclosure by any person as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all rights to the copyright whatsoever.

BACKGROUND OF THE INVENTION**Field of the Invention**

The invention relates generally to managing flows for a reservation protocol. More particularly, the invention relates to a technique for pre-allocating an aggregated reservation protocol session and thereafter sharing the reservation protocol session among multiple individual application sessions by multiplexing the multiple individual application flows thereon. The present invention also relates to management of a Voice Over Internet Protocol (VoIP) network via a novel Graphical User Interface (GUI) that enables a system manager to initialize, based on predicted link utilization, a plurality of routers and media aggregation managers existing on a selected communication path. The initialization provides the media aggregation managers with reservation protocol session parameters and bandwidth allocation requirements for a predetermined schedule of usage over the VoIP network.

Description of the Related Art

The transfer of voice traffic over packet networks, e.g., voice over Internet Protocol (VoIP), is rapidly gaining acceptance. However, significant work remains in the area of enhancing the quality of such services. One potential technique for improving the quality of voice calls involves the use of a bandwidth reservation protocol to communicate per-flow requirements by signaling the network. Typically, however, bandwidth reservation protocols require per-flow state information to be maintained at each intermediate node between the initiator and the prospective recipient. As a result, in a VoIP network relying on such bandwidth reservation protocols, scalability becomes an issue since each VoIP call reservation requires a non-trivial amount of ongoing message exchange, computation, and memory resources in each intervening node to establish and maintain the reservation.

An example of a particular bandwidth reservation protocol that illustrates this scalability problem is the Resource Reservation Protocol (RSVP). RSVP is an Internet Protocol- (IP) based protocol that allows applications running on end-stations, such as desktop computers, to communicate per-flow requirements by signaling the network. Using RSVP, the initiator of a VoIP call transmits a Path message downstream to the prospective recipient. The Path message causes state information, such as information regarding the reverse path to the initiator, to be stored in each node along the way. Subsequently, the prospective recipient of the VoIP call initiates resource reservation setup by communicating its requirements to an adjacent router via an upstream Resv message. For example, the prospective recipient may communicate a desired quality of service (QoS), e.g., peak/average bandwidth and delay bounds, and a description of the data flow to all intervening routers between the call participants. Additionally, after the reservation has been established, participating routers must continue to exchange periodic status and control messages to maintain the reservation. Consequently, processing and storage overhead associated with reservation establishment and maintenance increases linearly as a function of the number of calls. For further background and information regarding RSVP see Braden, R., Zhang, L., Berson, S., Herzog, S. and Jamin, S., "Resource Reservation Protocol (RSVP) Version 1 Functional Specification," RFC 2205, Proposed Standard, September 1997.

A proposed solution to RSVP's scalability problems can be found in F. Baker et al., "Aggregation of RSVP for IPv4 and IPv6 Reservations," Internet Draft, March 2000. However, the proposed solution requires a modification to RSVP, which would result in changes to router software. Additionally, the proposal does not use RSVP end-to-end, but rather uses Diff-Serv in the core. It may also require changes to other routing protocols like OSPF and IS-IS. Finally, it appears that there may also be additional burdens on network administrators to make the aggregation scheme work.

In light of the foregoing, what is needed is a less invasive technique for managing application flows that require real-time response, such as flows associated with VoIP services, and addressing scalability issues associated with bandwidth reservation protocols. It would also be desirable to minimize changes to the particular bandwidth reservation protocol employed and existing router software.

Current network management tools such as Hewlett Packard's Open View and AdventNet, have typically been used by System Administrators for detecting and analyzing faults that occur within a network. The programs generally discover a network and each node or router on the network submits to the administrator if and where faults exist in the network so that the System Administrator can address the problematic faults. The System administrator can select an individual router and provision the router through the Open View and AdventNet GUI. When provisioning a router, the existing tools utilize a standard protocol such as Simple Network Management Protocol (SNMP) or command line interface. The standard protocol is typically communicated to the provisioning tool like Open View or Adventnet by the router during the network discovery so that the protocol utilized for provisioning the router is hidden from the user. Provisioning a router includes router control parameters such as assigning an IP address to a router or assigning a bandwidth for a certain type of communication through the router.

BRIEF SUMMARY OF THE INVENTION

Apparatus and methods are described for multiplexing application flows over a pre-allocated bandwidth reservation protocol session and administering a VoIP network. According to one embodiment, a pre-allocated reservation protocol session is shared by one or more individual application sessions. The reservation protocol session is pre-allocated over a path between a first network device associated with a first user community and a second network device associated with a second user community based upon an estimated usage of the path for individual application sessions between users of the first and second user communities. Subsequently, the one or more individual application sessions are dynamically aggregated by multiplexing application flows associated with the one or more individual application sessions onto the pre-allocated reservation protocol session at the first network device and demultiplexing at the second network device.

According to a second embodiment, a network device enables multiple applications to share an aggregated reservation protocol session. The network device includes a storage device having stored therein one or more routines for establishing and managing the aggregated reservation protocol session. A processor coupled to the storage device executes the one or more routines to pre-allocate the aggregated reservation protocol session and thereafter share the aggregated reservation protocol session among multiple application sessions of individual application sessions. The aggregated reservation protocol session is pre-allocated based upon an estimate of the bandwidth requirements to accommodate the multiple application sessions. The aggregated reservation protocol session is shared by multiplexing, onto the aggregated reservation protocol session, outbound media packets originated by local application/endpoints associated with the application sessions, and demultiplexing, from the aggregated reservation protocol session, inbound media packets originated by remote application/endpoints.

According to a third embodiment, a method of conveying information about a VoIP network to a user is disclosed. The method comprises: discovering a plurality of nodes on a VoIP network wherein the plurality of nodes includes a media aggregation manager that provides application/protocol specific multiplexing/demultiplexing of media traffic onto a pre-allocated reservation protocol session; and graphically depicting representations of the plurality of nodes and their interconnections on a network map, wherein the

representations of the plurality of media aggregation managers are visually distinguishable from the remainder of the plurality of nodes.

According to a fourth embodiment, a method of allowing a user to interactively explore how changes in path selection between media aggregation managers affects projected link utilization in a network is disclosed. A graphical user interface displays graphical representations of a first media aggregation manager and second media aggregation manager. The first and second media aggregation managers serve as reservation session aggregation points between a first user community and a second user community and have a plurality of physical paths through which media packets may be exchanged by way of one or more packet forwarding devices. The GUI displays a first projected link utilization based upon an indication that a first path of the plurality of physical paths will be used to convey media packets between the first and second media aggregation managers. The GUI also displays a second projected link utilization based upon an indication that a second path of the plurality of physical paths will be used to convey media packets between the first and second media aggregation managers.

According to a fifth embodiment, a method is disclosed wherein the method comprises: in response to a discovery request, discovering nodes on a network; identifying and graphically displaying the nodes and their interconnections on a map; receiving inputs including a first node, a second node and projected bandwidth traffic requirement between the first node and the second node; and displaying the projected bandwidth traffic requirement for the nodes.

According to a sixth embodiment, a graphical user interface is disclosed wherein the GUI comprises: a display portion that graphically depicts and identifies a plurality of nodes on a network, wherein the plurality of nodes includes a plurality of media aggregation managers that provide application/protocol specific multiplexing/demultiplexing of media traffic onto a pre-allocated reservation protocol session, and wherein the plurality of media aggregation managers are distinguishable from other nodes on the network.

According to a seventh embodiment, a method is disclosed wherein the method comprises: receiving a first input indicating a first media aggregation manager; receiving a second input indicating a second media aggregation manager; receiving a third input indicating a projected utilization between the first media aggregation manager and the

second media aggregation manager; displaying a prioritized plurality of paths between the first media aggregation manager and the second media aggregation manager that satisfy the projected utilization; and receiving a fourth input indicating a selected path of the plurality of paths.

Other features of the present invention will be apparent from the accompanying drawings and from the detailed description that follows.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS.

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

Figure 1 conceptually illustrates interactions between two media aggregation managers according to one embodiment of the present invention.

Figure 2 is an example of a network device in which one embodiment of the present invention may be implemented.

Figure 3 is a high-level block diagram of a media aggregation manager according to one embodiment of the present invention.

Figure 4 is a simplified, high-level flow diagram illustrating media aggregation processing according to one embodiment of the present invention.

Figure 5 is a simplified, high-level flow diagram illustrating application session establishment processing according to one embodiment of the present invention.

Figure 6 illustrates interactions among local and remote media aggregation manager functional units according to one embodiment of the present invention.

Figure 7 is a flow diagram illustrating Registration, Admission, Status (RAS) signaling processing according to one embodiment of the present invention.

Figure 8 is a flow diagram illustrating call signaling processing according to one embodiment of the present invention.

Figure 9 is a flow diagram illustrating control signaling processing according to one embodiment of the present invention.

Figure 10 is a flow diagram illustrating media/control transmission multiplexing processing according to one embodiment of the present invention.

Figure 11 is a flow diagram illustrating media/control reception demultiplexing processing according to one embodiment of the present invention.

Figure 12 conceptually illustrates application session establishment in an H.323 environment according to one embodiment of the present invention.

Figure 13A illustrates the encapsulated ("MUX") packet format according to one embodiment of the present invention in which address replacement is performed by the LMAM.

Figure 13B illustrates media transmission in both directions according to the encapsulated packet format of **Figure 13A**.

Figure 14A illustrates the encapsulated ("MUX") packet format according to another embodiment of the present invention in which address replacement is performed by the RMAM.

Figure 14B illustrates media transmission in both directions according to the encapsulated packet format of **Figure 14A**.

Figure 15 illustrates an initialization control GUI in communication with a plurality of media aggregation Managers according to one embodiment of the present invention.

Figure 16 is a menu of available screens for the initialization GUI according to one embodiment of the present invention.

Figure 17 is a flow diagram illustrating a typical user navigation flow through the initialization process according to one embodiment of the present invention.

Figures 18 is a screen used for de-allocation of the media aggregation managers according to one embodiment of the present invention.

Figure 19 illustrates a network map interface according to one embodiment of the present invention.

Figure 20 illustrates a property window associated with a node according to one embodiment of the present invention.

Figure 21 illustrates a bandwidth allocation screen according to one embodiment of the present invention.

Figure 22 illustrates a BW on Link screen showing a utilization schedule for a selected node on the discovered network according to one embodiment of the present invention.

Figure 23 is a flow chart indicating the process of analysis for a selected path according to one embodiment of the present invention.

Figure 24 is a flow chart indicating the process of initializing the selected media aggregation managers according to one embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Apparatus and methods are described for multiplexing application flows over a pre-allocated bandwidth reservation protocol session and initializing, allocating and de-allocating reservation protocol sessions between a plurality of media aggregation managers. Broadly stated, embodiments of the present invention seek to provide a scalable architecture that enables efficient provisioning of reserved bandwidth for multiple application flows by multiplexing individual application flows over a pre-allocated reservation protocol session. Additionally, embodiments of the present invention seek to provide a graphical user interface (GUI) that enables a user to allocate and de-allocate bandwidth and reservation protocol sessions between a plurality of media aggregation managers along a path containing a plurality of routers. The pre-allocated reservation protocol session preferably takes into consideration current network resources and estimated usage of network resources, such as bandwidth, based upon historical data. For example, the amount of pre-allocated resources may vary due to different loads being offered at different times of day and/or day of week. Additionally, the pre-allocated reservation protocol session may be dynamically adjusted to account for actual usage that surpasses the estimated usage or actual usage that falls below the estimated usage.

Allocation and de-allocation of bandwidth and reservation protocol sessions between a plurality of media aggregation managers along a path containing a plurality of routers is facilitated by allowing the user to analyze various repercussions of increasing/decreasing the user demand over various paths on a Voice over Internet Protocol (VoIP) network and viewing the bandwidth effects at all nodes on the path for a

schedule that varies based on usage variations at various times of the day, week, month or year.

According to one embodiment, a more intelligent approach is employed in connection with initiation and maintenance of a large number of reservations. Rather than establishing and maintaining a reservation protocol session for each application flow that requires real-time response, which results in many independent reservation protocol sessions and high overhead, a single reservation protocol session may be pre-allocated and subsequently dynamically shared among the application flows by aggregating the associated media packets and transmitting them over a multiplexed media stream. For example, VoIP services may be provided between many different user communities using pre-allocated RSVP sessions between pairs of distributed media aggregation managers. The media aggregation managers multiplex outbound voice packets onto the pre-allocated RSVP session and demultiplex inbound voice packet received over the pre-allocated RSVP session, thereby sharing a common RSVP session and reducing the computational resources required by the network to provide real-time response for multiple application flows. Advantageously, in this manner, it becomes feasible to use reservation protocols, such as RSVP, for large numbers of applications that require real-time performance, such as VoIP services.

One benefit of the graphical user interface of the present invention is that it allows a system administrator to adjust bandwidth allocation requirements for a plurality of users communicating between a plurality of locations based on historical and current utilization demands by allowing allocation and de-allocation of bandwidth reservations between a plurality of media aggregation managers. Additionally, another advantage of the present invention is that the GUI allows a user, by selecting a path, to initialize multiple routers along the path simultaneously without having to individually provision each router. The present invention addresses the inadequacy of current network management tools by providing a GUI for discovering a VoIP network, including the media aggregations managers residing on the VoIP network and allowing a user, based on predicted usage requirements, to initialize the media aggregation managers and the routers included on a selected path for a predetermined schedule.

According to another embodiment, a VoIP network containing a plurality of media aggregation managers is discovered and then displayed. The user may review individual

properties for each of the nodes displayed on a network map. For example, the user may select two media aggregation managers for inter-communication analysis along with a predicted community demand of resources between the two selected media aggregation managers. The GUI displays a prioritized list of potential paths between the selected media aggregation managers including one or more routers for the communities to use in communicating between the media aggregation managers. Additionally, the user may select a path for an analysis of the effect of allocating the predicted bandwidth to a reservation protocol session between the selected media aggregation managers. The graphical user interface displays a predicted schedule of bandwidth traffic for any node on the network incorporating the predicted pre-allocated bandwidth that is being considered for allocation between the media aggregation managers. Based on the displayed data, the user may decide to allocate the bandwidth for all of the routers and media aggregation managers along the selected path, change paths, de-allocate bandwidth between these or other media aggregation managers or reduce/restrict the predicted community usage on a selected path.

In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention may be practiced without some of these specific details. In other instances, well-known structures and devices are shown in block diagram form.

The present invention includes various steps, which will be described below. The steps of the present invention may be performed by hardware components or may be embodied in machine-executable instructions, which may be used to cause a general-purpose or special-purpose processor programmed with the instructions to perform the steps. Alternatively, the steps may be performed by a combination of hardware and software.

The present invention may be provided as a computer program product which may include a machine-readable medium having stored thereon instructions which may be used to program a computer (or other electronic devices) to perform a process according to the present invention. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, magnet or optical cards, flash memory, or other type of media /

machine-readable medium suitable for storing electronic instructions. Moreover, the present invention may also be downloaded as a computer program product, wherein the program may be transferred from a remote computer to a requesting computer by way of data signals embodied in a carrier wave or other propagation medium via a communication link (e.g., a modem or network connection).

While, for convenience, embodiments of the present invention are described with reference to particular existing signaling, control, and communications protocol standards, such as International Telecommunication Union Telecommunication Standardization Section (ITU-T) Recommendation H.225.0 entitled "Call Signalling Protocols and Media Stream Packetization for Packet-based Multimedia Communication Systems," published February 1998 (hereinafter H.225.0); ITU-T Recommendation H.245 entitled "Control Protocol for Multimedia Communication," published May 1999 (hereinafter H.245); ITU-T Recommendation H.323 entitled "Packet-based Multimedia Communications Systems," published September 1999 (hereinafter H.323); and a particular bandwidth reservation protocol (i.e., RSVP), the present invention is equally applicable to various other signaling, control, communications and reservation protocols. For example, Session Initiation Protocol (SIP) may be employed to create, modify, and terminate application sessions with one or more participants. SIP is described in M. Handley et al., "SIP: Session Initiation Protocol," RFC 2543, Network Working Group, March 1999, which is hereby incorporated by reference.

In addition, for sake of brevity, embodiments of the present invention are described with reference to a specific application (i.e., VoIP) in which individual flows may be multiplexed over a pre-allocated bandwidth reservation protocol session. Nevertheless, the present invention is equally applicable to various other applications that require real-time performance, such as applications based on human interactions (e.g., collaborative software, online/Web collaboration, voice conferencing, and video conferencing), and the like.

Terminology

Brief definitions of terms used throughout this application are given below.

In the context of the described embodiment, a "media aggregation manager" may generally be thought of as a network device, such as an edge device at the ingress/egress

edges of a user community, or a group of one or more software processes running on a network device that provides application/protocol specific multiplexing/demultiplexing of media traffic onto a pre-allocated reservation protocol session.

A “reservation protocol” generally refers to a protocol that may be employed to communicate information regarding a desired level of service for a particular application flow. An example of an existing bandwidth reservation protocol is RSVP.

A “community” or “user community” generally refers to a group of users/residents residing on a common network at a given location. For example, employees on an enterprise network at a given location or users of a particular Internet service provider (ISP) at a given location may represent a user community.

In the context of the described embodiment, a “reservation protocol session” generally refers to a set of reserved network resources established and maintained between two or more network devices that serve as proxies for application endpoints residing behind the proxies. An example, of a reservation protocol session is an RSVP session between two media aggregation managers.

In the context of the described embodiment, an “application session” generally refers to a session established and maintained between two or more terminals. According to embodiments of the present invention, one or more application sessions may be multiplexed onto a single reservation protocol session thereby reducing the overhead for establishing and maintaining multiple application sessions.

“Total available bandwidth” refers to the amount of bandwidth accessible for any given router or could refer to the maximum available bandwidth of the most limiting node on a path between two selected nodes and their intervening nodes.

The “available communication bandwidth” encompasses the amount of bandwidth accessible for the desired type of communication to be reserved in any reservation protocol session. For instance, in one embodiment, the user may wish to allocate reservation protocol sessions for VoIP communication. In one case, 75% of the total available bandwidth may be the available communication bandwidth for VoIP type communications and a reservation protocol session initialized for 100 users between two media aggregation managers may only require 10% of the available communication bandwidth.

A “terminal” generally refers to a LAN-based endpoint for media transmission, such as voice transmission. Terminals may be capable of executing one or more

networked applications programs. An example of a terminal would be a computer system running an Internet telephony application, such as CoolTalk or NetMeeting.

An "application" or "endpoint" generally refers to a software program that is designed to assist in the performance of a specific task, such as Internet telephony, online collaboration, or video conferencing.

An "application flow" generally refers to the data associated with an application session. An example of an application flow is a media stream, such as a continuous sequence of packetized voice data transmitted over a network.

A "tag," in the context of the described embodiment, generally refers to information that is appended to application generated packets, such as Real-time Transport Protocol (RTP) packets or Real-time Transport Control Protocol (RTCP) packets, that allows the proxy endpoints of the reservation protocol session to transmit encapsulated packets to the appropriate remote application/endpoint (RA). According to one embodiment of the present invention, a tag includes address information, such as the destination network address of the terminal upon which the destination application/endpoint resides. When a media aggregation manager is employed in connection with a transport protocol and control protocol (such as RTP and RTCP) that use different channels or ports for control and data, control and data packets may be multiplexed onto the reservation protocol session as well by including protocol dependent control information. Then, the remote media aggregation manager may strip the tag from the encapsulated packet and determine the appropriate channel/port of the remote application/endpoint on which to forward the resulting packet based upon the additional protocol dependent control information within the tag. Advantageously, in this manner, two layers of multiplexing may be achieved, (1) a first layer that allows identification of the appropriate application at the remote media aggregation manager; and (2) a second layer that specifies a subclass/subprocess within an application.

Media Aggregation Overview

The architecture described herein seeks to resolve scalability problems observed in current reservation protocols. These scalability issues have slowed the adoption of reservation protocols in network environments where multiple applications must be provided with certainty regarding a minimum reserved bandwidth.

Figure 1 conceptually illustrates interactions between two media aggregation managers 110 and 120 according to one embodiment of the present invention. According to one embodiment, the media aggregation managers 110 and 120 act as reservation protocol proxies on behalf of the terminals 111, 112, 121, and 122. For example, the media aggregation managers 110 and 120 establish and maintain a reservation session, such as an RSVP session, between each other by exchanging reservation signaling messages 160. Subsequently, rather than establishing additional reservation protocol sessions, the media aggregation managers 110 and 120 respond to reservation requests from the terminals 111, 112, 121, and 122 by dynamically allocating the reserved resources, such as bandwidth, associated with the reservation protocol session to corresponding application sessions. In this manner, multiple application sessions may share the reservation session by multiplexing media packets onto the reservation session as described further below.

In this example, a multiplexed media/control stream 170 is established using admission control signaling messages 130. The multiplexed media/control stream 170 is carried over the pre-allocated reservation session between media aggregation manager 110 and media aggregation manager 120. The multiplexed media/control stream 170 represents one way to handle certain transport and control protocol combinations, such as RTP and RTCP, that use different channels or ports for control and data. In alternative embodiments, the reservation protocol session 160 may not need to distinguish between control and data.

While in the described embodiment, the media aggregation managers 110 and 120 are discussed as if they are autonomous network edge devices, it should be kept in mind that according to various embodiments of the present invention some or all of the functionality of a media aggregation manager might be integrated with existing network devices, such as bridges, routers, switches, and the like. Additionally, while only a single aggregated reservation protocol session between two media aggregation managers 110 and 120 is described in connection with the present example, it should be appreciated that each media aggregation manager 110 and 120 may support multiple, heterogeneous reservation protocol sessions capable of providing heterogeneous application flows among multiple user communities. Importantly, according to embodiments of the present invention, regardless of the number of terminals or application/endpoints, application flows may be

provided with reserved bandwidth between any and all pairs of terminals of N user communities by establishing and sharing no more than N^2 reservation protocol sessions.

Network Device Overview

An exemplary machine in the form of a network device 200, representing an exemplary media aggregation manager 110, in which features of the present invention may be implemented will now be described with reference to **Figure 2**. In this simplified example, the network device 200 comprises a bus or other communication means 201 for communicating information, and a processing means such as one or more processors 202 coupled with bus 201 for processing information. Networking device 200 further comprises a random access memory (RAM) or other dynamic storage device 204 (referred to as main memory), coupled to bus 201 for storing information and instructions to be executed by processor(s) 202. Main memory 204 also may be used for storing temporary variables or other intermediate information during execution of instructions by processor(s) 202. Network device 200 also comprises a read only memory (ROM) and/or other static storage device 206 coupled to bus 201 for storing static information and instructions for processor 202. Optionally, a data storage device (not shown), such as a magnetic disk or optical disc and its corresponding drive, may also be coupled to bus 201 for storing information and instructions.

One or more communication ports 225 may also be coupled to bus 201 for allowing various local terminals, remote terminals and/or other network devices to exchange information with the network device 200 by way of a Local Area Network (LAN), Wide Area Network (WAN), Metropolitan Area Network (MAN), the Internet, or the public switched telephone network (PSTN), for example. The communication ports 225 may include various combinations of well-known interfaces, such as one or more modems to provide dial up capability, one or more 10/100 Ethernet ports, one or more Gigabit Ethernet ports (fiber and/or copper), or other well-known interfaces, such as Asynchronous Transfer Mode (ATM) ports and other interfaces commonly used in existing LAN, WAN, MAN network environments. In any event, in this manner, the network device 200 may be coupled to a number of other network devices, clients and/or servers via a conventional network infrastructure, such as a company's Intranet and/or the Internet, for example.

Media Aggregation Manager

Figure 3 is a high-level block diagram of a media aggregation manager according to one embodiment of the present invention. By interconnecting a plurality of distributed media aggregation managers, such as media aggregation manager 300, an architecture is provided for multiplexing several application flows (e.g., VoIP calls) over a pre-allocated reservation protocol session, such as a pre-allocated RSVP pipe. Advantageously, the multiplexing of application flows reduces the computational resources required by the network to provide reserved bandwidth, e.g., guaranteed bandwidth, for multiple application flows. The source media aggregation manager receives media packets from its local terminals and transmits multiplexed media to the destination aggregation manager. The destination aggregation manager receives the multiplexed media and routes media packets to the appropriate terminal(s) of its local terminals.

In this example, the media aggregation manager 300 includes an application/protocol specific media multiplexor 350, an application/protocol specific media demultiplexor 360, an admission control manager 315, a generic resource manager 340, and a resource pool 345. In a software implementation, instances of the media multiplexor 350, media demultiplexor 360, and admission control manager 315 may be created for each particular application/protocol needed to allow communications between terminals of the geographically diverse user communities. Importantly, it should be appreciated that the particular partitioning of functionality described with reference to this example is merely illustrative of one or many possible allocations of functionality.

According to the embodiment depicted, the resource manager 340 establishes and maintains one or more pre-allocated reservation protocol sessions between the local media aggregation manager and one or more remote media aggregation managers. The resource manager 340 optionally interfaces with a centralized entity (not shown) that provides information relating to the characteristics and estimated amount of resources for the pre-allocated reservation protocol sessions. Alternatively, a network administrator may provide information to the resource manager 340 relating to desired characteristics of the pre-allocated reservation protocol sessions. The resource manager 340 also tracks active application sessions for each reservation protocol session and the current availability of resources for each reservation protocol session in the resource pool 345.

The admission control manager 315 interfaces with local terminals (not shown) associated with a particular user community, the media multiplexor 350, the resource manager 340, and one or more other remote media aggregation managers associated with other user communities. Importantly, in one embodiment, the media multiplexor 350 hides the details of how reserved resources are internally allocated and managed, thereby allowing the local terminals to use existing reservation protocols, such as RSVP, without change. The media multiplexor 350 receives media packets from the local terminals and appropriately translates/encapsulates the packets in accordance with the aggregation technique described further below. When application flows are established and terminated, the admission control manager 315 interfaces with the resource manager 340 to allocate and deallocate resources, respectively.

The media demultiplexor 360 interfaces with the local terminals to supply with media packets by demultiplexing their respective application flows from the pre-allocated reservation protocol session.

The admission control manager 315 exchanges admission control signaling messages with remote admission control managers and configures the local application/endpoint (LA) to send media to the media multiplexor 350 after an application session has been established with a remote media aggregation manager. For VoIP using the H.323 protocol, the admission control manager 315 may include RAS, call control, and call signaling processing.

In operation, two resource managers cooperate to establish a pre-allocated reservation protocol session between a local media aggregation manager (LMAM) and a remote media aggregation manager (RMAM). The resource managers make a reservation that is large enough to accommodate the anticipated load offered by applications that need to communicate over the reservation protocol session. Subsequently, a local media multiplexor (LMM) associated with the LMAM provides admission control for application flows between one or more terminals of the LMAM and the RMAM with the assistance of the local and remote admission control managers and the local and remote resource managers. If sufficient resources, such as bandwidth, are available over the pre-allocated reservation protocol session, then the local media multiplexor multiplexes the application flows over the pre-allocated reservation protocol session. On the receiving end, the remote media demultiplexor (RMD) demultiplexes the application flows and sends them to their

intended destinations. The typical admission control manager 315 will be a player in the path of the application protocol for setting up the connection between two or more application endpoints; hence, it may be instrumented to modify the path of the media packets to flow through the LMM and remote media multiplexor (RMM).

In brief, after an application session has been associated with the pre-allocated reservation protocol session, the application/endpoints may use a transport protocol and/or a control protocol, such as RTP and/or RTCP to exchange media packets between them. The media packets may carry various types of real-time data, such as voice, video, multi-media, or other data for human interactions or collaboration. Media packets from a data source are tagged by the local media multiplexor 350 and sent over the reserved path to one or more media demultiplexors 360 corresponding to the data destination. As illustrated below, the media demultiplexor 360 strips the tag before the media packets are forwarded and uses the tag information to determine the eventual destination of the data packet.

From the perspective of the local terminals, they are establishing and using reservation protocol sessions for each application flow. However, in reality, the media aggregation manager 300 shares the pre-allocated reservation protocol session among multiple application flows.

As will be described further below, a specific example of the use of this architecture is in connection with the use of the H.323 protocol for VoIP calls. Typically, an H.323 Gatekeeper is used by endpoints to help in address resolution, admission control etc. So, for the H.323 protocol, the gatekeeper is a convenient place for the media multiplexor 350 to reside.

Note that in this description, in order to facilitate explanation, the media aggregation manager 300 is generally discussed as if it is a single, independent network device or part of single network device. However, it is contemplated that the media aggregation manager 300 may actually comprise multiple physical and/or logical devices connected in a distributed architecture; and the various functions performed may actually be distributed among multiple network devices. Additionally, in alternative embodiments, the functions performed by the media aggregation manager 300 may be consolidated and/or distributed differently than as described. For example, any function can be

implemented on any number of machines or on a single machine. Also, any process may be divided across multiple machines.

Sharing a Pre-Allocated Reservation Protocol Session

Figure 4 is a simplified, high-level flow diagram illustrating media aggregation processing according to one embodiment of the present invention. In one embodiment, the processing blocks described below may be performed under the control of a programmed processor, such as processor 202. However, in alternative embodiments, the processing blocks may be fully or partially implemented by any programmable or hard-coded logic, such as Field Programmable Gate Arrays (FPGAs), TTL logic, or Application Specific Integrated Circuits (ASICs), for example.

In this example, it is assumed that, prior to the start of the media aggregation processing, a reservation protocol session has been established. The pre-allocated reservation protocol session preferably takes into consideration current network resources and estimated usage of network resources, such as bandwidth, based upon historical data. For example, the amount of pre-allocated resources may vary due to different loads being offered at different times of day and/or day of week.

At any rate, at decision block 410, the media aggregation manager 300 determines the type of event that has occurred. If the event represents the receipt of an application session establishment request from a local terminal, then processing proceeds to decision block 420. If the event represents the receipt of media packets from a local application/endpoint, then processing continues with decision block 450. If the event represents the receipt of a media packet from a remote application/endpoint, then control passes to processing block 460. If the event represents the receipt of an application session termination request, then processing continues with processing block 470.

At decision block 420, a determination is made whether resources are available to meet the needs identified in the application session establishment request. For example, the resource manager 340 may determine if sufficient bandwidth is available on an appropriate pre-allocated reservation protocol session by comparing a minimum bandwidth specified in the application session establishment request to a bandwidth availability indication provided by the resource pool 345.

If adequate resources are available to provide the requestor with the minimum resources requested, processing continues with processing block 430 where application session establishment processing is performed. Application session establishment processing is described below with reference to **Figure 5**. Otherwise, if there are insufficient resources to accommodate the application session establishment request, processing branches to processing block 440. At processing block 440, the media aggregation manager 300 may reject the application session establishment request. Alternatively, the media aggregation manager 300 may continue the application session establishment process and provide a best effort service for the request (without the use of pre-allocated resources).

At processing block 450, media packets received from a local application/endpoint are tagged and sent over the network to the destination using the previously reserved resources (e.g., the pre-allocated reservation protocol session). The tagging and multiplexing of media packets onto the pre-allocated reservation protocol session will be discussed in detail below.

At processing block 460, media packets received from a remote application/endpoint are forwarded to the appropriate local application/endpoint. For example, the packets may be sent to the appropriate local application/endpoint based upon an examination of the tag information added by the remote media aggregation manager.

At processing block 470, in response to an application session termination request, resources allocated to this application session are relinquished and made available for other application sessions. For example, the resource manager 340 may update an indication of available resources in the resource pool 345 for the pre-allocated reservation protocol session associated with the terminated application session.

Figure 5 is a simplified, high-level flow diagram illustrating application session establishment processing according to one embodiment of the present invention. In the present example, application session establishment processing begins with processing block 510. At processing block 510, the requested resources are allocated to the application session. According to one embodiment, the local resource manager 340 creates a new application session entry, in the resource pool 345, containing an indication of the resources granted to the application session.

At decision block 520, a determination is made whether the desired remote application/endpoint is available to participate in the application session. If so, processing proceeds to processing block 530; otherwise, processing branches to processing block 560.

Assuming the desired remote application/endpoint is available to participate in the application session, then at processing block 530, the local application/endpoint and the remote application/endpoint are configured to send media packets associated with the application session to the local and remote media multiplexors, respectively.

At processing blocks 540 and 550, the local and remote media multiplexors and demultiplexors are configured in accordance with the application session. For example, as described further below, a lookup table may be maintained by the media multiplexor 350 or media demultiplexor 360 to translate the source network address of the local application/endpoint to the destination network address of the remote application/endpoint.

Figure 6 illustrates interactions among local and remote media aggregation manager functional units according to one embodiment of the present invention. In general, the media aggregation managers abstract the true application session endpoints from each other and serve as proxies for their respective local applications/endpoints. The media aggregation managers accomplish this by intercepting messages originating from their respective local applications/endpoints and modifying the messages to make themselves appear as the actual application flow originators/recipients.

In this example, for simplicity, it is assumed that a single local application/endpoint (LA) is establishing an application session with a single remote application/endpoint (RA) over a pre-allocated reservation protocol session 690 between a local media aggregation manager (LMAM) geographically proximate to the LA and a remote media aggregation manager (RMAM) geographically proximate to the RA.

The LA transmits a request to connect to the RA to the LMAM (670). The LACM inquires of the local resource manager (LRM) whether sufficient resources are currently available to accommodate the LA's request (672). The LRM indicates the availability or inavailability of available resources to the LACM (674).

Assuming, sufficient resources are available to provide the reserved resources the LA needs for the requested connection to the RA, then the LACM asks the RACM if the RA is available (676). In response to the LACM's request, the RACM queries the RA to determine its present availability (678). The RA indicates whether or not it is currently available to participate in an application session (680).

Assuming, the RA indicates that it is available, then the RACM communicates the RA's availability to the LACM (682). In response to the availability of the RA, the LACM directs the RACM to proceed with establishment of a connection between the LA and RA.

Having determined that a connection is feasible, the LACM and RACM proceed to configure their media multiplexors and media demultiplexors for the LA-RA connection. The LACM configures the local media multiplexor (LMM) to tag media originated from the LA for routing to the RA and to send the resulting encapsulated media packets to the remote media demultiplexor (RMD) (686). The LACM further configures the local media demultiplexor (LMD) to forward media packets that are received from the RMM and tagged as being associated with the LA-RA connection to the LA (690).

Similarly, the RACM configures the remote media demultiplexor (RMD) to forward media packets that are received from the LMM and tagged as being associated with the LA-RA connection to the RA (688). The RACM also configures the remote media multiplexor (RMM) to tag media originated from the RA for routing to the LA and to send the resulting encapsulated media packets to the local media demultiplexor (LMD) (692).

Once the media multiplexors and media demultiplexors have been appropriately configured for the LA-RA connection, the LACM and the RACM inform their application/endpoints to commence transmission of media to the LMM and the RMM, respectively. Thus, the media aggregation managers appear to their respective application/endpoints as the actual application flow originators/recipients and subsequently serve as proxies for their respective application/endpoints.

During media transmission between the LA and the RA 698 and 699, media packets originated by the LA are sent to the LMM, which encapsulates the media packets by appending a tag appropriate for the LA-RA connection and forwards the encapsulated packets over the pre-allocated reservation protocol session 690 to the RMD. The RMD determines the RA is the intended destination based upon the tag, removes the tag, and forwards the media packet to the RA. Media packets originated by the RA are sent to the RMM, which encapsulates the media packets by appending a tag appropriate for the LA-RA connection and forwards the encapsulated packets over the pre-allocated reservation protocol session 690 to the LMD. The LMD determines the LA is the intended destination based upon the tag, removes the tag, and forwards the media packet to the LA.

An Exemplary H.323 VoIP Implementation

H.323 is basically an umbrella that covers several existing protocols, including but not limited to H.225.0, and H.245. The later two protocols are used to establish call connection, and capability information between two endpoints. Once this information is exchanged, the endpoints may use RTP and RTCP to exchange voice (and multi-media) information between them.

H.323 suggests that RTP/RTCP should be established between two endpoints (caller/receiver) for each call. Consequently, in order to provide Quality Of Service (QoS) for each call using a protocol like RSVP would mean that every endpoint pair

(caller/receiver) for every H.323 call would need to establish RSVP between one another. This would create a huge amount of overhead on the endpoint and adversely affect network resources as RSVP “soft states” must be maintained for the life of the call. This quickly becomes a tremendous scalability issue, since as number of simultaneous calls increase, so do the RSVP “soft state” maintenance messages between endpoints, and every router involved in the transmitting RTP/RTCP data stream.

The media aggregation manager 300 described herein seeks to provide a clean, and scalable solution for this problem, while providing the same QoS as if two individual endpoints had used a reservation protocol session, such as RSVP, between them. Briefly, according to the described H.323 VoIP embodiment, the H.323 endpoints (callers/receivers) need not have knowledge of how to establish and maintain RSVP sessions. Instead, the media aggregation managers establish one or more RSVP “pipes” between them that can accommodate several (expected) voice calls. These RSVP pipes are created as the media aggregation managers are started and the RSVP pipes are maintained between them. This immediately reduces the amount of RSVP state processing in the network. The RSVP pipes between media aggregation managers may be created based upon an educated estimate of the number of calls that are expected between user communities being managed by these media aggregation managers. Since RSVP by nature is established between a specific IP address/port pair and since the pipes are pre-created between media aggregation managers, all voice traffic (RTP/RTCP) originates and terminates between media aggregation managers at the media multiplexor 350 and the media demultiplexor 360, respectively.

In this manner, according to one embodiment, the “local” media aggregation manager appears to an H.323 voice application caller as its intended receiver. The H.323 endpoints make calls to the media multiplexors of the local media aggregation managers without realizing the local media aggregation managers are not really the final destination. The local media aggregation manager calls the remote media aggregation manager and passes the RTP/RTCP voice data to it. The remote media aggregation manager receives the voice data and sends it the “real” receiver while hiding all multiplexing details from both the caller and the receiver. However, as the voice data is actually exchanged between media aggregation managers over the network it gets RSVP treatment, reserved bandwidth, and QoS. Advantageously, this solution serves as a surrogate to route calls over the pre-

created RSVP pipes eliminating QoS processing by endpoints, without any deviations from each involved standard protocol.

Referring now to **Figure 7**, a flow diagram illustrating exemplary Registration, Admission, Status (RAS) signaling processing will now be described. At decision block 710, the appropriate processing path is determined based upon the triggering event. If the event is a request for a terminal's signaling address then processing proceeds to decision block 720. If the event represents a signaling address response, then control flow branches to processing block 750. However, if the event is a new call request, then processing continues with decision block 760.

At decision block 720, in response to a request for a terminal signaling address, a determination is made whether or not the terminal is locally serviced. If it is determined that the terminal is not serviced by the media aggregation manager 300, then processing continues with processing block 730; otherwise processing proceeds to processing block 740.

At processing block 730, the media aggregation manager 300 requests the call signaling address from an appropriate remote media aggregation manager. For example, the local media aggregation manager may transmit a multicast message or a directed broadcast to locate the appropriate remote media aggregation manager that services the desired terminal.

At processing block 740, the media aggregation manager 300 returns its own signaling address rather than the signaling address of the locally serviced terminal. In this manner, subsequent call signaling and control signaling is routed through the local media aggregation manager rather than letting the locally service terminal handle this signaling directly.

At processing block 750, in response to a signaling address response, the media aggregation manager 300, as above, returns its signaling address in place of the signaling address of the locally serviced terminal to abstract call and control signaling from the locally serviced terminal.

At decision block 760, in response to a new call request on the RAS channel of the media aggregation manager 300, a determination is made whether there is capacity for the new call. For example, the local resource manager verifies whether the reservation protocol session over which the new call will be multiplexed can accommodate the

additional bandwidth requirements of the new call. At any rate, if the local resource manager determines that the reservation protocol session has adequate resources for the new call, then processing continues to processing block 770. Otherwise, control flows to processing block 780.

At processing block 770, the media aggregation manager 300 returns an indication that the new call can be accepted. At processing block 780, the media aggregation manager 300 returns direction to reject the new call.

Figure 8 is a flow diagram illustrating call signaling processing according to one embodiment of the present invention. At decision block 810, the appropriate processing path is determined based upon the event that has triggered the call signaling processing tread. If the event is a local call connect request, the processing proceeds to processing block 820. If the event represents a remote call connect request, then control flow branches to processing block 830. If the event is a local alerting/call or proceeding/connect message, then processing continues with processing block 840. However, if the event is a remote alerting/call or proceeding/connect message, the processing proceeds with processing block 850.

At processing block 820, in response to a local call connect request, the media aggregation manager 300 accepts the call from the local terminal and calls the remote media aggregation manager that services the destination terminal. In this manner, the local media aggregation manager poses as the intended receiver to its local terminals that are callers.

At processing block 830, in response to a remote call connect request, the media aggregation manager 300 accepts the call from the remote media aggregation manager and calls the intended recipient, e.g., one of the terminals serviced by the local media aggregation manager. In this manner, the local media aggregation manager poses a caller to its local terminals that are receivers.

At processing block 840, in response to a local alerting/call or proceeding/connect message, the local media aggregation manager relays the message to the appropriate remote media aggregation manager(s).

At processing block 850, in response to a remote alerting/call or proceeding/connect message, the local media aggregation manager relays the message to

the appropriate local terminal(s). After processing block 850, call signaling is complete and control protocol signaling (e.g., H.245) can begin.

Figure 9 is a flow diagram illustrating control signaling processing according to one embodiment of the present invention. At decision block 910, the appropriate processing path is determined based upon the event that has triggered the control signaling processing tread. If the event is receipt of a master/slave and capability exchange from a local application/endpoint, the processing proceeds to processing block 920. If the event represents receipt of a master/slave and capability exchange from a remote media aggregation manager, then control flow branches to processing block 930. If the event is receipt of logical channel information from a local application/endpoint, then processing continues with processing block 940. However, if the event is reception of logical channel information from a remote media aggregation manager, the processing proceeds with processing block 950.

At processing block 920, the master/slave and capability exchange is transmitted to the remote media aggregation manager.

At processing block 930, the master/slave and capability exchange is transmitted to the local application/endpoint.

At processing block 940, the logical channel information from the local application/endpoint is stored in anticipation of making a connection with the media and/or control channels of the local application/endpoint. At processing block 950, the LMAM forwards its own logical channel information to the RMAM. Additionally, the network address of the LA is sent to the RMAM.

At processing block 960, the network address of the RA is stored in a lookup table for address translation and the logical channel information of the LMAM is forwarded to the LA.

Figure 10 is a flow diagram illustrating media/control transmission multiplexing processing according to one embodiment of the present invention. At processing block 1010, the local media multiplexor reports the resources being consumed by the local application/endpoint to the local resource manager.

At processing block 1020, the media aggregation manager 300 connects to the media and/or control channels of the local application/endpoint.

At processing block 1030, media and control data packets are generated by the local application/endpoint and received by the local media multiplexor. The media multiplexor 350 takes packets coming from either the control or media channels of the local application/endpoint and sends them to the appropriate remote media aggregation manager(s). According to this example, the media multiplexor 350 marks the outbound packets with appropriate address information (referred to as a “tag”) for demultiplexing at the remote media aggregation manager. The tag is typically appended to transport protocol packets, such as TCP or RTP packets, to allow the media multiplexor 350 to direct packets to the appropriate remote application/endpoint. According to one embodiment, the tag includes address information, such as the destination network address associated with the remote application/endpoint. The destination network address may be determined with reference to a lookup table that allows translation between the source network address associated with the local application/endpoint and the destination network address associated with the remote application/endpoint. Alternatively, a lookup table may be maintained on the media demultiplexor 360 and the tag would include the source network address associated with the local application/endpoint. Then, the source network address would be used by the remote media demultiplexor to determine how to route the inbound packet to the appropriate remote application/endpoint.

When different channels or ports are used for transport and control protocols (such as RTP and RTCP), then the tag may also include additional protocol dependent control information to allow multiplexing of data and control packets onto the reservation protocol session. Therefore, at optional processing block 1050, each outbound packet may additionally be marked as control or data to allow the remote media aggregation manager to determine the appropriate channel/port of the remote application/endpoint on which to forward the packet.

Finally, at processing block 1060, the marked packet is transmitted to the appropriate remote media aggregation manager(s).

Figure 11 is a flow diagram illustrating media/control reception demultiplexing processing according to one embodiment of the present invention. At processing block

1110, a packet is received from a remote media aggregation manager. The demultiplexing information (e.g., the tag) added by the remote media multiplexor is stripped from the packet and examined at processing block 1120. Optionally, at processing block 1130, if control and data packets are being multiplexed onto the reservation protocol session, a determination is made whether the packet is a media packet or a control packet based upon the tag. At processing block 1140, the appropriate local application(s)/endpoint(s) to which the packet is destined is/are determined. As described above, the media multiplexor 350 may perform address translation from a source network address to a destination network address. In this case, the media demultiplexor 360 determines the appropriate local application(s)/endpoint(s) that are to receive the packet by examining the address portion of the tag. Alternatively, if the media multiplexor 350 leaves the source network address in the address portion of the tag, then the media demultiplexor 360 determines the appropriate local application(s)/endpoint(s) by first translating the address portion using a local lookup table, for example.

In any event, finally, at processing block 1150, the media demultiplexor 360 transmits the packet to those of the local application(s)/endpoint(s) identified in processing block 1140. If, according to the particular transport and/or control protocols employed, the application(s)/endpoint(s) receive media packets and control packets on different channels/ports, then the packet is forwarded onto the appropriate channel/port of the local application(s)/endpoints(s) based on the packet classification performed at processing block 1130.

Figure 12 conceptually illustrates application session establishment in an H.323 environment according to one embodiment of the present invention. In general, the media aggregation managers abstract the true application session endpoints from each other and serve as proxies for their respective local applications/endpoints. As explained above, the media aggregation managers accomplish this by intercepting signaling messages originating from their respective local applications/endpoints and modifying the signaling messages to make themselves appear as the actual callers/recipients.

In this illustration, for simplicity, it is assumed that a single local application/endpoint (LA) is establishing an application session with a single remote application/endpoint (RA) over a pre-allocated reservation protocol session 1290 between

a local media aggregation manager (LMAM) geographically proximate to the LA and a remote media aggregation manager (RMAM) geographically proximate to the RA.

According to this example, application session establishment involves RAS signaling 1210 and 1230, H.225 signaling 1240, and H.245 signaling 1250. RAS signaling 1210 begins with a request for the RA signaling address 1211 by the LA to the LMAM. The LMAM transmits the request 1211 via the reservation protocol session 1290 to the RMAM. In response to the request 1211, the RMAM decides it wants to route H.225/H.245 signaling through it instead of letting the RA do it directly. Therefore, the RMAM replies with a packet 1212 containing RMAM's signaling address. Similarly, the LMAM decides it wants to route H.225/H.245 signaling through it instead of letting the LA do it directly. Therefore, the LMAM substitutes its signaling address for that of the RMAM and forwards packet 1213 to the LA.

RAS signaling continues with the RA asking the RMAM (on its RAS channel) if it is okay to accept a new call by sending the RMAM a new call request 1231. The RMAM authorizes the new call by responding with a packet 1231 giving the RA permission to accept the new call.

H.225 signaling comprises the RA sending H.225 alerting/call proceeding/connect messages 1241 to the RMAM. The RMAM sends the same to the LMAM; and the LMAM sends the same to the LA. At this point, the LA determines that H.225 call signaling is complete and starts H.245 signaling.

H.245 signaling begins with the LA sending master/slave and capability exchange messages 1251 to the LMAM, which are relayed to the RMAM and from the RMAM to the RA. Then, the RA sends master/slave and capability exchange messages 1252 to the RMAM. The RMAM transmits these messages to the LMAM; and the LMAM forwards them to the LA.

Subsequently, the LA initiates an exchange of logical channel information by sending logical channel information packets 1253 to the LMAM. The logical channel information identifies the network address (e.g., IP address) and port numbers where RTP/RTCP connections will be accepted. The LMAM stores the LA's logical channel information and passes its own connection information 1254 to the RMAM. Additionally, the LMAM provides the network address of the LA to the RMAM for later use in address translation lookups. As mentioned above, the network address of the LA may be used by

the RMM or the RMD depending upon where the address translation lookup is performed. The RMAM remembers the information provided by the LMAM and generates its own RTP/RTCP information 1255 and passes it to the RA.

After receiving logical channel information thought to be associated with the LA, the RA sends its logical channel information 1256 to the RMAM (thinking it is being directed to the LA). The RMAM stores the RA's logical channel information and passes its own connection information 1257 to the LMAM. Additionally, the RMAM provides the network address of the RA to the LMAM. The LMAM remembers the logical channel information provided by the RMAM and generates its own RTP/RTCP information 1258 and passes it to the LA.

The LA sends an ACK message 1259 to the LMAM to acknowledge receipt of what it thinks to be the RA's logical channel information. The acknowledgement is relayed to the RA by the LMAM and the RMAM. The RA also sends an ACK message 1260 to the RMAM to acknowledge receipt of what it thinks to be the LA's logical channel information. The acknowledgement is relayed to the LA by the RMAM and the LMAM. Finally, the LMAM and the RMAM each use the logical channel information intercepted from the LA and the RA, respectively, to connect to the media and/or control channels of the LA and RA.

Exemplary Encapsulated Packet Formats

Figure 13A illustrates the encapsulated ("MUX") packet format 1300 according to one embodiment of the present invention in which address replacement is performed by the LMAM. The payload of the encapsulated packet 1300 includes a destination network address field 1310, a variable length transport or control protocol packet portion 1315, and a packet type indication 1320. The destination network address 1310 is typically the IP address of the true recipient (e.g., the application/endpoint to which the packet is destined). In environments where multiplexing of control and data is employed, the variable length portion 1315 may include either a transport protocol packet (e.g., a RTP packet) or a control protocol packet (e.g., a RTCP packet) as indicated by the packet type indication 1320. In alternative embodiments, where multiplexing of control and data is not employed, then the variable length portion 1315 would still include either control or data, but the packet type indication 1320 would no longer be necessary.

Figure 13B illustrates media transmission in both directions according to the encapsulated packet format of **Figure 13A**. When the LA originates a media packet, it generates a packet 1340 including media 1342. The LMAM encapsulates the media 1342 in the encapsulated packet format 1300 by generating an encapsulated packet 1350 that includes the RA's network address 1351, the media 1342, and a packet type indicator 1353. For example, upon receipt of packet 1340, the LMAM may append the network address of the RA and a packet type indicator 1353 based upon the channel/port upon which the packet 1340 was received. When the encapsulated packet 1350 is received by the RMAM, it strips the information added by the LMAM and forwards a packet 1360 comprising the media 1342 to the RA.

When the RA originates a media packet, it generates a packet 1390 including media 1392. The RMAM encapsulates the media 1392 in the encapsulated packet format 1300 by generating an encapsulated packet 1380 that includes the LA's network address 1341, the media 1392, and a packet type indicator 1383. For example, upon receipt of packet 1390, the RMAM may append the network address of the LA and a packet type indicator 1383 based upon the channel/port upon which the packet 1390 was received. When the encapsulated packet 1380 is received by the LMAM, it strips the information added by the RMAM and forwards a packet 1370 comprising the media 1392 to the LA.

Figure 14A illustrates the encapsulated ("MUX") packet format according to another embodiment of the present invention in which address replacement is performed by the RMAM. The payload of the encapsulated packet 1400 includes a source network address field 1410, a variable length transport or control protocol packet portion 1415, and a packet type indication 1420. The source network address 1410 is typically the IP address of the true caller (e.g., the application/endpoint from which the packet is originated). In environments where multiplexing of control and data is employed, the variable length portion 1415 may include either a transport protocol packet (e.g., a RTP packet) or a control protocol packet (e.g., a RTCP packet) as indicated by the packet type indication 1420. In alternative embodiments, where multiplexing of control and data is not employed, then the variable length portion 1415 would still include either control or data, but the packet type indication 1420 would no longer be necessary.

Figure 14B illustrates media transmission in both directions according to the encapsulated packet format of **Figure 14A**. When the LA originates a media packet, it generates a packet 1440 including media 1442. The LMAM encapsulates the media 1442 in the encapsulated packet format 1400 by generating an encapsulated packet 1450 that includes the LA's network address 1441, the media 1442, and a packet type indicator 1453. For example, upon receipt of packet 1440, the LMAM may append the network address of the LA and a packet type indicator 1453 based upon the channel/port upon which the packet 1440 was received. When the encapsulated packet 1450 is received by the RMAM, it strips the information added by the LMAM and forwards a packet 1460 comprising the media 1442 to the RA by looking up the network address of the RA based upon the LA's network address 1441.

When the RA originates a media packet, it generates a packet 1490 including media 1492. The RMAM encapsulates the media 1492 in the encapsulated packet format 1400 by generating an encapsulated packet 1480 that includes the RA's network address 1451, the media 1492, and a packet type indicator 1483. For example, upon receipt of packet 1480, the RMAM may append the network address of the RA and a packet type indicator 1483 based upon the channel/port upon which the packet 1480 was received. When the encapsulated packet 1480 is received by the LMAM, it strips the information added by the RMAM and forwards a packet 1470 comprising the media 1492 to the RA by looking up the network address of the LA based upon the RA's network address 1451.

Media Aggregation Manager Overview

Figure 15 conceptually illustrates interactions between two media aggregation managers 1530 and 1540 according to one embodiment of the present invention. The media aggregation managers 1530 and 1540 act as reservation protocol proxies on behalf of the communities 1550 and 1560 where a plurality of residents wish to communicate with each other. For example, resident 1551 may wish to communicate with resident 1561 while resident 1552 wishes to communicate with resident 1562. The media aggregation managers pre-allocate bandwidth and establish a reservation protocol session capable of handling multiple communications between residents in Community 1550 and residents in Community 1560. Having media aggregation managers controlling a single reservation protocol session for multiple communication for residents between a plurality of

communities allows for packets of communication data to be efficiently multiplexed and reduces protocol overhead as individual pairs of residents need not maintain their own application sessions.

The reservations may apply to various paths. For example, the bandwidth reservation may lay over path 1510 containing one intermediary router 1511 or may be allocated over path 1520 containing two intermediary routers 1521 and 1522. The reservation for communications between community 1550 and community 1560 may also be split over the various paths 1510 and 1520 depending on the historical and current bandwidth burden on individual routers 1511, 1521 and 1522. The media aggregation managers reserve a protocol session and then multiplex the plurality of data packets for a plurality of communication links to be communicated. As prior technologies required each resident in a community to request an individual reservation session to establish a link between Community 1550 and Community 1560, media aggregation managers and the apparatuses and methods required for initializing/controlling the media aggregation managers have been developed. The present invention focuses on the graphical user interface 1500 that enables a user to interactively discover, analyze and initialize the media aggregation managers to handle a schedule of community communications.

The administration GUI tool used for initializing the routers and media aggregation managers is illustrated as designator 1500 in Figure 15. The instructions for the GUI may reside in any combination of hardware or software and likewise may reside on any system configured to interact with other nodes on the network.

Graphical User Interface Overview

Figure 16 demonstrates one embodiment of a navigation tool for accessing various screens of the graphical user interface. In the embodiment depicted, a user may choose from one of the listed options, for instance, a user may select Network Discovery 1601 to discover the network to be initialized or may choose Bandwidth Allocation 1603 to allocate bandwidth to or establish a reservation protocol session between selected media aggregation managers as will become apparent in the following description.

An example of how a user may navigate through the menu to administer to a network is depicted in **Figure 17**. Beginning with the menu depicted in Figure 16, a user may select Network Discovery 1601 in processing block 1710. Once the Network

Discovery 1710 is complete, the user may select to display the network map by selecting Network Map 1602 from the menu. After viewing the network map that displays all or a subset of the communities, nodes and media aggregation managers currently on the system, the user may choose to go directly to the Bandwidth Allocation screen 1603 by selecting the menu link or may choose to right-click on a graphical representation of one of the media aggregation managers and select from a pop-up menu to allocate bandwidth for that particular media aggregation manager. In either case, a Bandwidth Allocation screen presents itself to the user enabling him to select two media aggregation managers and indicate the number of users capable of communicating via the selected media aggregation managers 1730. Once the user indicates which media aggregation managers are to be allocated and how many users are predicted to utilize the session, one or more potential paths between the two media aggregation managers are displayed on the bandwidth allocation interface. The user may select a path for analysis and, through the graphical user interface, indicate that the selected path is to be analyzed. At processing block 1740, the selected path is analyzed to determine projected bandwidth utilization for each link of the selected path. Once analyzed, a user may select BW on Link 1606 from the menu or the BW on Link screen may automatically appear after analysis has completed.

On the BW on Link screen, the user may select any node on the network, specifically of interest would be those altered by the predicted increase in usage. In response to being selected, the screen displays a schedule of usage for that node and optionally a projection indicating if the predicted usage increase is within an acceptable range 1750. When the predicted usage is within an acceptable range, the media aggregation managers may be initialized. In one embodiment, the user selects Bandwidth Allocation 1603 from the menu and, based on the nodes all falling within an acceptable range, the bandwidth for the selected media aggregation managers 1760 and the routers along the path is allocated. The user can then decide if more media aggregation managers need to be allocated 1770 (for instance, if a pre-existing plurality of communities are experiencing an increase of residents in the near future). When no more media aggregation managers need to be initialized, then the initialization is complete 1780. On the other hand, when more media aggregation managers need to be initialized, the user may return to the network map interface through the Network Map menu item 1602 or may return

directly to the Bandwidth Allocation Interface through the Bandwidth Allocation menu item 1603 and repeat the media aggregation selection process just described.

Alternatively, if the BW on Link screen provides data indicating that the predicted bandwidth utilization on any portion of the schedule exceeds the limitations of the network, the user may choose to select a different path for analysis or select to de-allocate a previously allocated session between two other media aggregation managers 1790. In either case, the user may return to the Bandwidth allocation page to select a different path through the bandwidth allocation menu item 1603 or the user may select a different combination of media aggregation managers to analyze or de-allocate. If the user decides to de-allocate a session between two selected media aggregation managers to make available more bandwidth to accomplish the desired decrease in predicted utilization, the user may simply select the media aggregation managers 1800 and then click on the menu option Bandwidth Deallocation 1604 which brings up a dialog box 1820 and de-allocate screen 1830, shown in **Figure 18**, allowing the user to de-allocate the current session between the selected media aggregation managers 1810.

Network Map Interface

Figure 19 shows the network map interface according to one embodiment of the invention. A graphical representation of a plurality of nodes on the discovered network is shown. In addition, links between each of the nodes and the administration GUI 1950 are shown. The network map screen indicates community nodes 1910, router nodes 1920 and media aggregation managers 1930. Each of the nodes or media aggregation nodes are visually distinct via a graphical representation indicative of the type of node. The user is able to readily identify whether a node is a community, router, media aggregation manager, & etc. simply by looking at its graphical representation. The community nodes 1910 may have a plurality of residents, including but not limited to computers, routers, phones, printers, scanners and the like. Each of the nodes and the media aggregation managers have properties associated with it that may be accessed by positioning the cursor over the graphical representation for the node and clicking on a mouse button assigned for property retrieval, in this embodiment, although not shown, the right mouse button is assigned for property retrieval. A properties window immediately appears as shown in **Figure 20** indicating information about the node such as the manufacturer 2010, the interface

addresses 2020 or a name 2030. Additionally, the properties window may indicate other information about the characteristics of the current configuration of the node. For instance, the property window for a media aggregation manager may indicate how many reservation protocol sessions it is maintaining and with which other media aggregation managers each of the reservation protocol sessions are concerning. The property window may also indicate the available bandwidth for a given node and for what type of communication the bandwidth is available, such as voice or data communication and the amount of bandwidth that is currently allocated for reservation protocol sessions utilizing this particular media aggregation manager as a proxy or gate-keeper. Other properties may include interface command options, such as allocate bandwidth 1940, de-allocate bandwidth (not shown), or other interface command options that take the user to various interface screens and option windows.

Figure 21 is a snapshot of one embodiment of the bandwidth allocation screen. The user may select two community gate-keepers or media aggregation managers 2110 for analysis or initialization. The present embodiment allows the user to select a source media aggregation manager 2120, in this case “reddog” from a menu listing all media aggregation managers that were discovered on the network (not shown) and a destination media aggregation manager 2130, in this case “rossini”. The user may also designate the number of users 2140 capable of communicating from each of the selected media aggregation managers. In this example, 100 users are capable of simultaneously communicating through the media aggregation manager reddog to residents whose gate-keeper or media aggregation manager is Rossini and likewise, 100 users are capable of communicating from Rossini to residents of reddog. Although the number of users for this example is 100 for both media aggregation managers, they need not be the same number of users.

Once the user has selected two media aggregation managers for analysis or initialization, the user may select “OK” 2150 to indicate to the graphical user interface’s processing algorithms to evaluate all available paths between the two media aggregation managers. The user may also decide to “abort” the path evaluation process by selecting the “abort” button 2160.

In this example, two paths are determined during the path evaluation process although the invention is not so limited. The graphical user interface then displays the paths graphically depicting all intervening communities, routers or other nodes that lie

between the selected media aggregation managers. The graphical user interface may display the list in a prioritized fashion utilizing factors such as the number of nodes between the media aggregation managers, the physical length of travel between nodes, the total available bandwidth on the nodes between the media aggregation managers, the available communication bandwidth, or the propagation speed between the various nodes that make up the path. For each factor or combination of weighted factors, the most limiting of the intervening nodes may be utilized for the computation as would be readily apparent to one skilled in the art.

The user may then select a path 2170 to analyze. In most cases, the user may default to the highest prioritized path that in this case defaults to the first position on the graphical user interface but may be configured by the user to appear where desired. Alternatively, the user may see that a node in the prioritized path is going to ultimately be extremely burdened by other allocations that the user needs to initialize or has already been initialized and instead may opt for a lower prioritized path. In either case, according to this example, once the user has selected a path for allocation or analysis, he then chooses whether to reserve the protocol session between the two media aggregation managers by pressing the “start bandwidth Allocation” button 2180 or the user may select to analyze the effect the bandwidth allocation would have on the nodes by selecting “analyze selected path” button 2190.

The bandwidth allocation screen allows the user to abort the analysis at any time if so desired by selecting the “abort” button 2160.

Figure 23 demonstrates what happens when the analyze button 2190 is selected. In step 2310, a schedule of bandwidth allocation is determined for the selected path. In step 2320, after the predicted schedule for the selected path has been determined, the schedule of increased bandwidth allocation is overlaid on top of the schedule that accounts for bandwidth previously reserved to the nodes on the path via other media aggregation managers utilizing those nodes. Finally, in step 2340, the combined schedule of usage is optionally displayed to the user.

Once the analysis of the selected path has completed, the graphical user interface may automatically switch to the BW on Link screen shown in **Figure 22** or the user may select BW on Link from the menu on the left and previously discussed with regard to **Figure 16**. The BW on Link screen, in this embodiment, displays the predicted utilization

results of reserving the session as indicated on the Bandwidth allocation interface. As previously indicated, the displayed schedule incorporates all previously allocated sessions and bandwidth reservations burdening the intervening nodes as well as the predicted increase as a result of the analyzed path if it were to be allocated. The results of the analysis may be viewed for each of the nodes displayed in the network map, primarily of interest would be the nodes along the selected path so that a determination can be made as to whether the protocol session to be reserved will exceed the available communication bandwidth for any node at any time in the predicted schedule.

The media aggregation managers that have been analyzed are displayed 2210. The user may indicate a time range for display by changing the offset for each router 2220. Another segment of the display 2230 indicates to the user all available and analyzed nodes between the selected media aggregation managers by way of a scrollable list of intervening nodes. The user may then select a node on the path and a schedule of utilization for that node appears 2240. The schedule. The schedule depicts a time frame including a Start Time 2250 and End time 2260 and indicates the bandwidth utilized during that time frame 2270 and the amount of the available communication bandwidth 2280 that would remain available after the analyzed path has been allocated. The schedule covers various segments of the day as determined by the offsets selected 2220 and also indicates a schedule of usage for the node for various days of the week. Once the user verifies that the utilization on all of the nodes on the path are within a desirable range, the user may select to return to the bandwidth allocation screen shown in **Figure 21** and allocate the bandwidth 2180.

Once the allocate bandwidth button 780 is selected, the bandwidth for the media aggregation managers are allocated as shown in the flow chart in **Figure 24**. In this example at step 2410, each and every router on the selected path where RSVP is not currently utilized, RSVP is enabled. In Step 2420, each router on the selected path is provisioned to force all communication media between the residents communicating between selected source and destination media aggregation managers to travel across the media aggregation managers and routers of the selected path. In step 2430, the media aggregation managers are initialized with all scheduling information necessary to reserve protocol sessions for the plurality of residents at any time within the schedule. The reservation protocol sessions manage the protocol sessions for multiple communication links in order to reduce the overhead and delay times occurring when individual links must

be maintained as in previous technologies. The necessary scheduling information may include information such as how much bandwidth needs to be allocated for each session, expected increases and decreases in utilization based on time and other information necessary to manage a reservation protocol session. In step 2440, the media aggregation managers begin reserving protocol sessions according to the information schedule provided in step 2430.

In some instances, for example where the schedule indicates that utilization will exceed the available communication bandwidth, the user may select another path for analysis, select another pair of allocated media aggregation nodes for de-allocation or restrict the number of users allowed to communicate over the selected media aggregation managers. Should the user decide to de-allocate a previously allocated protocol session, he selects the media aggregation managers and then selects Bandwidth Deallocation 1604 from the menu. **Figure 18** indicates a bandwidth deallocation screen and allows the user to select "deallocate bandwidth". In response, the graphical user interface provides a warning and confirmation dialog box. The user may then confirm the deallocation.

In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

CLAIMS

What is claimed is:

1. A method comprising:
pre-allocating a reservation protocol session over a path between a first network device associated with a first user community and a second network device associated with a second user community based upon an estimated usage of the path for individual application sessions between users of the first user community and the second user community; and
dynamically aggregating one or more individual application sessions by multiplexing application flows associated with the one or more individual application sessions onto the pre-allocated reservation protocol session at the first network device.
2. The method of claim 1, further comprising demultiplexing the application flows at the second network device.
3. The method of claim 1, wherein the reservation protocol session comprises a Resource Reservation Protocol (RSVP) session.
4. The method of claim 1, wherein at least one of the one or more individual application sessions comprises a H.323 session.
5. The method of claim 1, wherein the reservation protocol session comprises a Resource Reservation Protocol (RSVP) session (RSVP) session and at least one of the one or more individual application sessions comprises a H.323 session.
6. The method of claim 1, further comprising:
forming an encapsulated packet by appending a tag to a media packet received at the first network device from a source local application/endpoint associated with the first user community, the tag including information to allow the second network device to determine a destination local application/endpoint associated with the second user community; and

removing the tag at the second network device prior to forwarding the media packet to the destination local application/endpoint.

7. The method of claim 6, wherein the tag includes a network address associated with the source local application/endpoint.
8. The method of claim 6, wherein the tag includes a network address associated with the destination local application/endpoint.
9. The method of claim 6, wherein the tag includes a packet type indicator that specifies how to further identify a subprocess within the destination local application/endpoint.
10. A method comprising:
establishing an aggregated reservation protocol session over a path between a first edge device and a second edge device based upon an estimate of the number of individual application sessions needed for the path during a predetermined window of time; and
sharing the aggregated reservation protocol session among a plurality of individual application sessions by tagging packets associated with corresponding application flows for transmission between the first edge device and the second edge device, the tagged packets being multiplexed onto the aggregated reservation protocol session by the first edge device or the second edge device and demultiplexed by the other.
11. A method comprising:
pre-allocating an aggregated reservation protocol session over a path between a first media aggregation manager and second media aggregation manager of a plurality of distributed media aggregation managers based upon an estimate of the bandwidth requirements for individual application sessions over the path;
sharing the aggregated reservation protocol session among a plurality of individual application sessions established between applications running on devices residing behind the first and second media aggregation managers by tagging

packets at the source media aggregation manager of the first or second media aggregation manager that resides in front of the source of a particular transmission, multiplexing the tagged packets onto the aggregated reservation protocol session, and demultiplexing the tagged packets at the destination media aggregation manager of the first or second media aggregation manager that resides in front of the destination of the particular transmission.

12. A method comprising:
 - establishing a Resource Reservation Protocol (RSVP) session between a first network device and a second network device that are part of an Internet Protocol (IP) network;
 - receiving, at the first network device from a first local terminal, a request to initiate a first H.323 session with a first remote terminal associated with the second network device;
 - allocating a portion of pre-allocated resources associated with the RSVP session to the first H.323 session between the first local terminal and the first remote terminal;
 - receiving, at the first network device from a second local terminal, a request to initiate a second H.323 session with a second remote terminal associated with the second network device;
 - allocating a portion of the pre-allocated resources associated with the RSVP session to the second H.323 session between the second local terminal and the second remote terminal; and
 - sharing the RSVP session between the first H.323 session and the second H.323 session by multiplexing voice packets associated with the first and second H.323 sessions onto the RSVP session.
13. The method of claim 12, further comprising:
 - transmitting voice packets from the first local terminal and first remote terminal by forming an encapsulated packet at the first network device that includes tag information to allow the second network device to determine the voice packets are intended for the first remote terminal; and

removing the tag information at the second network device prior to forwarding the voice packets to the first remote terminal.

14. The method of claim 13, wherein the tag information includes the IP address of the first local terminal.
15. The method of claim 13, wherein the tag information includes the IP address of the first remote terminal.
16. The method of claim 13, wherein the tag information includes a packet type indicator that specifies how to further identify a subprocess within the first remote terminal.
17. A method comprising:
establishing a Resource Reservation Protocol (RSVP) session between a first network device and a second network device that are part of an Internet Protocol (IP) network;
the first network device presenting itself as the recipient of a first call originated by a first local terminal associated with the first network device by providing its logical channel information to the first local terminal rather than providing logical channel information associated with the intended recipient of the first call;
the first network device presenting itself as the recipient of a second call originated by a second local terminal associated with the first network device by providing its logical channel information to the second local terminal rather than providing logical channel information associated with the intended recipient of the second call; and
the first network device transmitting voice packets from the first local terminal to the intended recipient of the first call and from the second local terminal to the intended recipient of the second call by multiplexing the voice packets onto the RSVP session.
18. The method of claim 17, further comprising the first network device presenting itself as the originator of a third call to the second network device by providing its

logical channel information to the second network device rather than providing logical channel information associated with a third local terminal associated with the first network device that truly originated the third call.

19. A media aggregation manager comprising:
 - a resource manager to establish a reservation protocol session with one or more other media aggregation managers prior to establishment of any application sessions that share resources associated with the reservation protocol and to subsequently allocate and deallocate the resources in response to anticipated use of application session establishment requests and application session termination requests, respectively;
 - an admission control manager coupled to the resource manager, the admission control manager to provide admission control for application flows based upon availability of the resources as indicated by the resource manager;
 - a media multiplexor coupled to the admission control manager, the media multiplexor to tag media packets received from local application/endpoints that are associated with admitted application flows and to transmit the tagged media packets over the reservation protocol session; and
 - a media demultiplexor to forward media packets received from remote application/endpoints to the local application/endpoints based upon tags appended by a media multiplexor of the one or more other media aggregation managers.
20. A network device comprising:
 - a storage device having stored therein one or more routines for establishing and managing an aggregated reservation protocol session;
 - a processor coupled to the storage device for executing the one or more routines to pre-allocate the aggregated reservation protocol session and thereafter share the aggregated reservation protocol session among a plurality of individual application sessions, where:
 - the aggregated reservation protocol session is pre-allocated based upon an estimate of the bandwidth requirements to accommodate the plurality of individual application sessions,

the plurality of individual application sessions are established between a plurality of local application/endpoints and a plurality of remote application/endpoints,
the aggregated reservation protocol session is shared by multiplexing outbound media packets originated by a local application/endpoint of the plurality of local application/endpoints onto the aggregated reservation protocol session, and demultiplexing inbound media packets originated by a remote application/endpoint of the plurality of remote application/endpoints from the aggregated reservation protocol session.

21. A system for multiplexing individual application sessions over a pre-allocated reservation protocol session comprising:
a first edge device coupled to a first plurality of terminals, the first edge device including a multiplexor to provide admission control for application sessions and to multiplex packets of admitted application flows over the pre-allocated reservation protocol session; and
a second edge device coupled to a second plurality of terminals, the second edge device including a demultiplexor to forward packets associated with the admitted application flows to the appropriate terminal of the second plurality of terminals.
22. A method of conveying information about a Voice Over Internet Protocol (VoIP) network to a user comprising:
discovering a plurality of nodes on the VoIP network, the plurality of nodes including a plurality of media aggregation managers that provide application/protocol specific multiplexing/demultiplexing of media traffic onto a preallocated reservation protocol session; and
graphically depicting representations of the plurality of nodes and their interconnections on a network map, wherein the representations of the plurality of media aggregation managers are visually distinguishable from the remainder of the plurality of nodes.

23. The method of claim 22, further comprising displaying a plurality of physical paths that are available for exchanging media packets between a selected pair of media aggregation managers of the plurality of media aggregation managers.
24. The method of claim 23, wherein the plurality of physical paths are prioritized in terms of their relative desirability for serving as the path over which media packets will be transferred between the first and second media aggregation managers.
25. A method of allowing a user to interactively explore how changes in path selection between media aggregation managers affects projected link utilization in a network comprising:
 - displaying graphical representations of a first media aggregation manager and a second media aggregation manager, the first and second media aggregation managers serving as reservation session aggregation points between a first user community and a second user community and having a plurality of physical paths through which media packets may be exchanged by way of one or more packet forwarding devices;
 - displaying a first projected link utilization based upon an indication that a first path of the plurality of physical paths will be used to convey media packets between the first and second media aggregation managers; and
 - displaying a second projected link utilization based upon an indication that a second path of the plurality of physical paths will be used to convey media packets between the first and second media aggregation managers.
26. The method of claim 23, further comprising overlaying a selected path of the plurality of physical paths onto existing bandwidth allocations to determine a projected link utilization associated with the selected path.
27. A method comprising:
 - in response to a discovery request, discovering nodes on a network;
 - identifying and graphically displaying the nodes and their interconnections on a map;

receiving inputs including a first node, a second node and a projected bandwidth traffic between the first node and the second node; and displaying a projected bandwidth utilization for the nodes that accounts for the increase in bandwidth utilization caused by the projected bandwidth traffic for a schedule.

28. The Method of claim 27 wherein the nodes include at least one media aggregation manager.
29. The method of claim 28 further comprising displaying a plurality of paths between the first node and the second node.
30. The method of claim 29 where the plurality of paths between the first node and the second node are prioritized by a criteria.
31. A Graphical User Interface (GUI) comprising:
a display portion that graphically depicts and identifies a plurality of nodes on a network, wherein the plurality of nodes includes a plurality of media aggregation managers that provide application/protocol specific multiplexing/demultiplexing of media traffic onto a preallocated reservation protocol session, and wherein the plurality of media aggregation managers are distinguishable from other nodes on the network.
32. The GUI of Claim 31 further comprising an identification table for displaying characteristics of a selected node.
33. A method utilizing a Graphical User Interface (GUI) comprising:
receiving a first input indicating a first media aggregation manager;
receiving a second input indicating a second media aggregation manager;
receiving a third input indicating a projected utilization between the first media aggregation manager and the second media aggregation manager;
displaying a prioritized plurality of paths between the first media aggregation manager and the second media aggregation manager that satisfy the projected utilization; and

receiving a fourth input indicating a selected path of the plurality of paths.

34. The method of Claim 33 further comprising a control initializing an allocation of bandwidth between the first media aggregation manager and the second media aggregation manager.
35. The method of claim 34 wherein the allocation of bandwidth comprises a provisioning of plurality of routers between the first media aggregation manager and the second media aggregation manager.
36. The method of claim 35 wherein the provisioning of the plurality of routers includes instructions that force media to route through the plurality of routers when being communicated from a first community of residents utilizing the first media aggregation manager to a second community of residents utilizing the second media aggregation manager.
37. The Method of Claim 36 further comprising an analysis control for receiving an input indicating the initiation of analysis of the first path.
38. The method of Claim 36 further comprising:
receiving a fifth input indicating a node on the selected path; and
displaying a schedule projecting bandwidth utilization for the node.
39. A method comprising substantially simultaneously provisioning a plurality of routers to force a media to travel from a first media aggregation manager through the plurality of routers and to a second media aggregation manager.
40. A method comprising provisioning a plurality of routers according to a path selected by a user over which reservation protocol session packets are forced to travel.
41. The method of claim 40 wherein the path includes an endpoint wherein the endpoint is a media aggregation manager.




Figure 1





Figure 2

**Figure 3**





Figure 4

**Figure 5**




Figure 6




Figure 7




Figure 8






Figure 9

**Figure 10**

**Figure 11**




Figure 12




Figure 13A




Figure 13B




Figure 14A




Figure 14B




Figure 15




Fig. 16




Figure 17




Fig. 18




Fig 19




Fig. 20




Fig. 21

Warning: Applet Window

Configured Communities

Source Job	Destination Job
reddog	rossini

DETAILS

Source	Destination
Name: reddog	Name: rossini
Time Offset(minutes): 120	Time Offset(minutes): -60
Connected Router IP: 192.168.13.250	Connected Router IP: 192.168.10.250

Links between the pair

Path Segment #	Link IP Address
1	192.168.13.250
2	192.168.12.250
3	192.168.12.252
4	192.168.1.2

Monday 2240

Start Time(NET5)	End Time(NET5)	B/W Used(bytes)
07:29:00	11:02:00	4000
10:29:00	14:02:00	4000
11:02:00	11:59:00	2000
11:59:00	16:59:00	5000
14:02:00	14:59:00	2000

Tuesday

Start Time(NET5)	End Time(NET5)	B/W Used(bytes)
07:29:00	11:02:00	4000
10:29:00	14:02:00	4000
11:02:00	11:59:00	2000
11:59:00	16:59:00	4000
14:02:00	14:59:00	2000

Wednesday

CONFIGURED COMMUNITIES

Fig. 22




Figure 23




Figure 24

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US01/24878

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) :Please See Extra Sheet.

US CL :Please See Extra Sheet.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : Please See Extra Sheet.

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y, P	US 6,298,120 B1 (CIVANLAR ET AL.) 02 OCTOBER 2001, COL.2-6	1-41
Y	US 6,009,469 A (MATTAWAY ET AL.) 28 Dec, 1999, COL.3-34.	1-41
Y, P	US 6,243,376 B1 (NG ET AL.) 05 JUNE 2001, COL.2-12.	1-41
A, P	US 6,243,759 B1 (BODEN ET AL.) 05 JUNE 2001, COL.2-8.	1-41
A, P	US 6,259,771 B1 (KREDO ET AL.) 10 JULY 2001, COL.1-4.	1-41
A, P	US 6,208,638 B1 (RIELEY ET AL.) 27 MARCH 2001, COL.2-6.	1-41

Further documents are listed in the continuation of Box C. See patent family annex.

"T"	Special categories of cited documents:	
"A"	document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E"	earlier document published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L"	document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O"	document referring to an oral disclosure, use, exhibition or other means	"Z" document member of the same patent family
"P"	document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

30 OCTOBER 2001

Date of mailing of the international search report

19 DEC 2001

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231
Facsimile No. (703) 305-8280

Authorized officer *Peggy Hanod*
GLEN BURGESS
Telephone No. (703) 306-4529

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US01/24878

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A, P	US 6,226,678 B1 (MATTAWAY ET AL) 01 MAY 2001, COL.2-30.	1-41

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US01/24878

A. CLASSIFICATION OF SUBJECT MATTER:
IPC (7):

G06F 18/00, 18/88, 17/00; H04L 12/66; H04M 1/64, 5/00, 7/00, 11/00

A. CLASSIFICATION OF SUBJECT MATTER:
US CL :

870/852, 853, 854, 880, 889, 892, 401; 879/88.17, 219, 220.01, 244, 901; 709/203, 217, 218, 219, 220, 221, 223, 224, 225, 227, 230, 238

B. FIELDS SEARCHED

Minimum documentation searched

Classification System: U.S.

870/852, 853, 854, 880, 889, 892, 401; 879/88.17, 219, 220.01, 244, 901; 709/203, 217, 218, 219, 220, 221, 223, 224, 225, 227, 230, 238